Optimization Letters manuscript No.
(will be inserted by the editor)

A Parallel Branch and Bound Algorithm for the Maximum Labelled Clique Problem

Ciaran McCreesh · Patrick Prosser

arXiv:1407.7061v2 [cs.DS] 17 Nov 2014

Received: date / Accepted: date

Abstract The maximum labelled clique problem is a variant of the maximum clique problem where edges in the graph are given labels, and we are not allowed to use more than a certain number of distinct labels in a solution. We introduce a new branch-and-bound algorithm for the problem, and explain how it may be parallelised. We evaluate an implementation on a set of benchmark instances, and show that it is consistently faster than previously published results, sometimes by four or five orders of magnitude.

Keywords Maximum labelled clique · Parallel branch and bound · Combinatorial optimisation · Computational experiments

1 Introduction

A clique in a graph is a set of vertices, where every vertex in this set is adjacent to every other in the set. Finding the size of a maximum clique in a given graph is one of the fundamental NP-hard problems. Carrabs et al. [1] introduced a variant called the maximum labelled clique problem. In this variant, each edge in the graph has a label, and we are given a budget b: we seek to find as large a clique as possible, but the edges in our selected clique may not use more than b different labels in total. In the case that there is more than one such maximum, we must find the one using fewest different labels. We illustrate these concepts in Figure 1, using an example graph due to Carrabs et al.; our four labels are shown using different styled edges.

This work was supported by the Engineering and Physical Sciences Research Council [grant number EP/K503058/1]

C. McCreesh and P. Prosser
University of Glasgow
Glasgow, Scotland
E-mail: c.mccreesh.1@research.gla.ac.uk and E-mail: patrick.prosser@glasgow.ac.uk

Carrabs et al. give example applications involving social network analysis and telecommunications. For social network analysis, vertices in the graph may represent people, and labelled edges describe some kind of relationship such as a shared interest. We are then seeking a large, mutually connected group of people, but using only a small number of common interests. For telecommunications, we may wish to locate mirroring servers in different data centres, all of which must be connected for redundancy. Labels here tell us which companies operate the connections between data centres: for simplicity and cost, we have a budget on how many different companies' connections we may use.

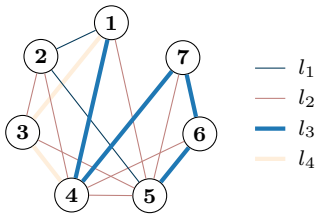

Fig. 1 A graph with maximum clique $\{1, 2, 3, 4, 5\}$, using all four edge labels. If our budget is only three, a maximum feasible clique has size four. There are several such cliques, but $\{4, 5, 6, 7\}$ is optimal since it uses only two labels, whilst every other uses at least three.

A mathematical programming approach to solving the problem was presented by Carrabs et al., who used CPLEX to provide experimental results on a range of graph instances. Here we introduce the first dedicated algorithm for the maximum labelled clique problem, and then describe how it may be parallelised to make better use of today's multi-core processors. We evaluate our implementation experimentally, and show that it is consistently faster than that of Carrabs et al., sometimes by four or five orders of magnitude. These results suggest that state of the art maximum clique algorithms are not entirely inflexible, and can sometimes be adapted to handle side constraints and a more complicated objective function without losing their performance characteristics.

Definitions and Notation Throughout, let $G = (V, E)$ be a graph with vertex set V and edge set E. Our graphs are undirected, and contain no loops. Associated with G is a set of labels, and we are given a mapping from edges to labels. We are also given a budget, which is a strictly positive integer.

The *neighbourhood* of a vertex is the set of vertices adjacent to it, and its *degree* is the cardinality of its neighbourhood. A *colouring* of a set of vertices is an assignment of colours to vertices, such that adjacent vertices are given different colours. A *clique* is a set of pairwise-adjacent vertices. The *cost* of a clique is the cardinality of the union of the labels associated with all of its edges. A clique is *feasible* if it has cost not greater than the budget. We say that a feasible clique C' is *better than* a feasible clique C if either it has larger cardinality, or if it has the same cardinality but lower cost. The *maximum labelled clique problem* is to find a feasible clique which is either better than or equal to any other feasible clique in a given graph—that is, of all the maximum feasible cliques, we seek the cheapest.

The hardness of the maximum clique problem immediately implies that the maximum labelled clique problem is also NP-hard. Carrabs et al. showed

that the problem remains hard even for complete graphs, where the maximum clique problem is trivial.

2 A Branch and Bound Algorithm

In Algorithm 1 we present the first dedicated algorithm for the maximum labelled clique problem. This is a branch and bound algorithm, using a greedy colouring for the bound. We start by discussing how the algorithm finds cliques, and then explain how labels and budgets are checked.

Algorithm 1: An algorithm for the maximum labelled clique problem.

```
1  maximumLabelledClique :: (Graph G, Int budget) → Vertex Set
2  begin
3      permute G so that vertices are in non-increasing degree order
4      global (C*, L*) ← (∅, ∅)
5      expand(true, ∅, every vertex of G, ∅)
6      expand(false, ∅, every vertex of G, ∅)
7      return C* (unpermuted)

8  expand :: (Boolean first, Vertex Set C, Vertex Set P, Label Set L)
9  begin
10     (order, bounds) ← colourOrder(P)
11     for i ← |P| downto 1 do
12         if |C| + bounds[i] < |C*| or (first and |C| + bounds[i] = |C*|) then
13             return
14         v ← order[i]
15         add v to C
16         L' ← L ∪ the labels of edges between v and any vertex in C
17         if |L'| ≤ (budget if first, otherwise |L*| − 1) then
18             if (C, L') is better than (C*, L*) then (C*, L*) ← (C, L')
19             P' ← the vertices in P that are adjacent to v
20             if P' ≠ ∅ then expand(first, C, P', L')
21         remove v from C and from P

22 colourOrder :: (Vertex Set P) → (Vertex Array, Int Array)
23 begin
24     (order, bounds) ← ([], [])
25     uncoloured ← P
26     colour ← 1
27     while uncoloured ≠ ∅ do
28         colourable ← uncoloured
29         while colourable ≠ ∅ do
30             v ← the first vertex of colourable
31             append v to order, and colour to bounds
32             remove v from uncoloured and from colourable
33             remove from colourable all vertices adjacent to v
34         add 1 to colour
35     return (order, bounds)
```

Branching: Let v be some vertex in our graph. Any clique either contains only v and possibly some vertices adjacent to v, or does not contain v. Thus we may build up potential solutions by recursively selecting a vertex, and branching on whether or not to include it. We store our growing clique in a variable C, and vertices which may potentially be added to C are stored in a variable P. Initially C is empty, and P contains every vertex (line 5).

The `expand` function is our main recursive procedure. Inside a loop (lines 11 to 21), we select a vertex v from P (line 14). First we consider including v in C (lines 15 to 20). We produce a new P' from P by rejecting any vertices which are not adjacent to v (line 19)—this is sufficient to ensure that P' contains only vertices adjacent to *every* vertex in C. If P' is not empty, we may potentially grow C further, and so we recurse (line 20). Having considered v being in the clique, we then reject v (line 21) and repeat.

Bounding: If we can colour a graph using k colours, we know that the graph cannot contain a clique of size greater than k (each vertex in a clique must be given a different colour). This gives us a bound on how much further C could grow, using only the vertices remaining in P. To make use of this bound, we keep track of the largest feasible solution we have found so far (called the *incumbent*), which we store in $C^\star$. Initially $C^\star$ is empty (line 4). Whenever we find a new feasible solution, we compare it with $C^\star$, and if it is larger, we unseat the incumbent (line 18).

For each recursive call, we produce a constructive colouring of the vertices in P (line 10), using the `colourOrder` function. This process produces an array *order* which contains a permutation of the vertices in P, and an array of bounds, *bounds*, in such a way that the subgraph induced by the first i vertices of *order* may be coloured using $bounds[i]$ colours. The *bounds* array is non-decreasing ($bounds[i+1] \geq bounds[i]$), so if we iterate over *order* from right to left, we can avoid having to produce a new colouring for each choice of v. We make use of the bound on line 12: if the size of the growing clique plus the number of colours used to colour the vertices remaining in P is not enough to unseat the incumbent, we abandon search and backtrack.

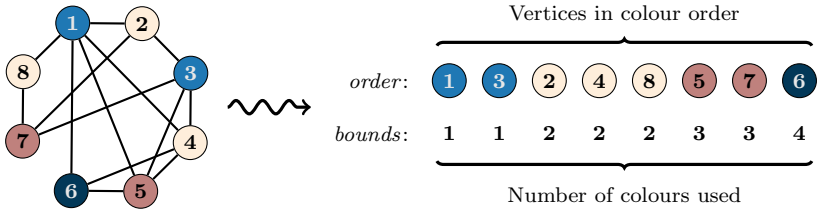

Fig. 2 The graph on the left has been coloured greedily: vertices 1 and 3 were given the first colour, then vertices 2, 4 then 8 were given the second colour, then vertices 5 and 7 were given the third colour, then vertex 6 was given the fourth colour. On the right, we show the *order* array, which contains the vertices in the order which they were coloured. Below, the *bounds* array, containing the number of colours used so far.

The `colourOrder` function performs a simple greedy colouring. We select a vertex (line 30) and give it the current colour (line 31). This process is repeated until no more vertices may be given the current colour without causing a conflict (lines 29 to 33). We then proceed with a new colour (line 34) until every vertex has been coloured (lines 27 to 34). Vertices are placed into the *order* array in the order in which they were coloured, and the ith entry of the *bounds* array contains the number of colours used at the time the ith vertex in *order* was coloured. This process is illustrated in Figure 2.

Initial vertex ordering: The order in which vertices are coloured can have a substantial effect upon the colouring produced. Here we will select vertices in a static non-increasing degree order. This is done by permuting the graph at the top of search (line 3), so vertices are simply coloured in numerical order. This assists with the bitset encoding, which we discuss below.

Labels and the budget: So far, what we have described is a variation of a series of maximum clique algorithms by Tomita et al. [13,12,14] (and we refer the reader to these papers to justify the vertex ordering and selection rules chosen). Now we discuss how to handle labels and budgets. We are optimising subject to two criteria, so we will take a two-pass approach to finding an optimal solution.

On the first pass (*first* = **true**, from line 5), we concentrate on finding the largest feasible clique, but do not worry about finding the cheapest such clique. To do so, we store the labels currently used in C in the variable L. When we add a vertex v to C, we create from L a new label set L' and add to it any additional labels used (line 16). Now we check whether we have exceeded the budget (line 17), and only proceed with this value of C if we have not. As well as storing $C^\star$, we also keep track of the labels it uses in $L^\star$.

On the second pass (*first* = **false**, from line 6), we already have the size of a maximum feasible clique in $|C^\star|$, and we seek to either reduce the cost $|L^\star|$, or prove that we cannot do so. Thus we repeat the search, starting with our existing values of $C^\star$ and $L^\star$, but instead of using the budget to filter labels on line 17, we use $|L^\star| - 1$ (which can become smaller as cheaper solutions are found). We must also change the bound condition slightly: rather than looking only for solutions strictly larger than $C^\star$, we are now looking for solutions with size equal to $C^\star$ (line 12). Finally, when potentially unseating the incumbent (line 18), we must check to see if either C is larger than $C^\star$, or it is the same size but cheaper.

This two-pass approach is used to avoid spending a long time trying to find a cheaper clique of size $|C^\star|$, only for this effort to be wasted when a larger clique is found. The additional filtering power from having found a clique containing only one additional vertex is often extremely beneficial. On the other hand, label-based filtering using $|L^\star| - 1$ rather than the budget is not possible until we are sure that $C^\star$ cannot grow further, since it could be that larger feasible maximum cliques have a higher cost.

Bit parallelism: For the maximum clique problem, San Segundo et al. [11,10] observed that using a bitset encoding for SIMD-like parallelism could speed up an implementation by a factor of between two to twenty, without changing the steps taken. We do the same here: P and L should be bitsets, and the graph should be represented using an adjacency bitset for each vertex (this representation may be created when G is permuted, on line 3). Most importantly, the *uncoloured* and *colourable* variables in `colourOrder` are also bitsets, and the filtering on line 33 is simply a bitwise and-with-complement operation.

Note that C should not be stored as a bitset, to speed up line 16. Instead, it should be an array. Adding a vertex to C on line 15 may be done by appending to the array, and when removing a vertex from C on line 21 we simply remove the last element—this works because C is used like a stack.

Thread parallelism: Thread parallelism for the maximum clique problem has been shown to be extremely beneficial [2,5]; we may use an approach previously described by the authors [5,7] here too. We view the recursive calls to `expand` as forming a tree, ignore left-to-right dependencies, and explore subtrees in parallel. For work splitting, we initially create subproblems by dividing the tree immediately below the root node (so each subproblem represents a case where $|C| = 1$ due to a different choice of vertex). Subproblems are placed onto a queue, and processed by threads in order. To improve balance, when the queue is empty and a thread becomes idle, work is then stolen from the remaining threads by resplitting the final subproblems at distance 2 from the root.

There is a single shared incumbent, which must be updated carefully. This may be stored using an atomic, to avoid locking. Care must be taken with updates to ensure that $C^\star$ and $L^\star$ are compared and updated simultaneously—this may be done by using a large unsigned integer, and allocating the higher order bits to $|C^\star|$ and the lower order bits to the bitwise complement of $|L^\star|$.

Note that we are not dividing a fixed amount of work between multiple threads, and so we should not necessarily expect a linear speedup. It is possible that we could get no speedup at all, due to threads exploring a portion of the search space which would be eliminated by the bound during a sequential run, or a speedup greater than the number of threads, due to a strong incumbent being found more quickly [3]. A further complication is that in the first pass, we could find an equally sized but more costly incumbent than we would find sequentially. Thus we cannot even guarantee that this will not cause a slowdown in certain cases [15].

3 Experimental Results

We now evaluate an implementation of our sequential and parallel algorithms experimentally. Our implementation was coded in C++, and for parallelism, C++11 native threads were used. The bitset encoding was used in both cases.

Table 1 Experimental results. For each graph, we use three different label set sizes and three different budgets, with randomly allocated labels, and show averages over 100 runs. In each case, we show the average size and cost of the result, the sequential runtime in seconds, the parallel runtime in seconds (2 cores, 4 threads) and then the "Enhanced" times reported by Carrabs et al. [1].

		25% budget					50% budget					75% budget				
Instance	$\|L\|$	size	cost	t_{seq}	t_{par}	[1]	size	cost	t_{seq}	t_{par}	[1]	size	cost	t_{seq}	t_{par}	[1]
johnson8-2-4	4	3.13	1.00	0.00	0.00	0.03	4.00	1.87	0.00	0.00	0.02	4.00	1.87	0.00	0.00	0.01
28 vertices,	8	3.51	1.50	0.00	0.00	0.03	4.00	2.48	0.00	0.00	0.01	4.00	2.48	0.00	0.00	0.02
210 edges	12	4.00	2.85	0.00	0.00	0.01	4.00	2.85	0.00	0.00	0.02	4.00	2.85	0.00	0.00	0.02
MANN_a9	11	5.64	2.76	0.01	0.00	3.18	8.89	5.93	0.02	0.01	13.98	13.34	8.99	0.01	0.00	6.47
45 vertices,	21	6.61	5.60	0.02	0.01	12.43	9.74	10.68	0.08	0.02	46.52	13.32	15.73	0.03	0.01	15.55
918 edges	32	7.00	7.79	0.04	0.01	17.61	10.26	14.99	0.19	0.05	108.13	14.12	23.28	0.03	0.01	13.18
hamming6-4	6	3.99	1.97	0.00	0.00	0.32	4.00	1.99	0.00	0.00	0.33	4.00	1.99	0.00	0.00	0.35
64 vertices,	11	4.00	2.64	0.00	0.00	0.43	4.00	2.64	0.00	0.00	0.42	4.00	2.64	0.00	0.00	0.37
704 edges	17	4.00	2.98	0.00	0.00	0.41	4.00	2.98	0.00	0.00	0.43	4.00	2.98	0.00	0.00	0.44
hamming6-2	15	6.07	3.82	0.04	0.01	41.32	9.87	7.86	0.76	0.21	382.73	15.42	12.00	2.44	0.66	711.20
64 vertices,	29	7.17	7.12	0.41	0.11	221.77	11.01	14.73	6.25	1.70	2368.78	15.86	21.84	11.73	3.14	2079.50
1824 edges	44	8.00	10.81	1.35	0.37	429.21	12.00	21.49	20.19	5.47	4955.96	17.13	32.65	23.15	6.28	2170.84
johnson8-4-4	14	5.99	3.93	0.01	0.00	49.82	7.98	6.94	0.02	0.01	85.73	11.09	10.89	0.01	0.00	16.36
70 vertices,	27	6.30	5.95	0.03	0.01	102.60	9.07	12.81	0.02	0.01	110.53	12.18	20.29	0.00	0.00	11.80
1855 edges	40	7.01	8.97	0.05	0.01	177.13	10.00	19.01	0.02	0.01	60.79	12.97	28.93	0.00	0.00	2.39
johnson16-2-4	23	6.50	5.28	0.15	0.04	4248.29	8.00	8.91	0.14	0.04	1943.33	8.00	8.91	0.14	0.04	2066.26
120 vertices,	46	7.75	11.17	0.29	0.08	5660.29	8.00	12.21	0.22	0.06	3044.75	8.00	12.21	0.22	0.06	3290.52
5460 edges	69	8.00	14.23	0.28	0.08	3699.71	8.00	14.23	0.28	0.08	4227.76	8.00	14.23	0.28	0.08	3909.15
keller4	28	6.98	6.89	0.28	0.07	$> 3h$	9.04	12.68	0.10	0.03	$> 3h$	11.00	18.75	0.02	0.01	3304.30
171 vertices,	55	8.00	12.85	0.32	0.09	$> 3h$	11.00	26.98	0.02	0.01	4081.97	11.00	26.98	0.02	0.01	4173.34
9435 edges	83	9.00	19.82	0.14	0.04	$> 3h$	11.00	31.88	0.02	0.01	4827.99	11.00	31.88	0.02	0.01	5028.16

Experimental results are produced on a desktop machine with an Intel i5-3570 CPU and 12GBytes of RAM. This is a dual core machine, with hyper-threading, so for parallel results we use four threads (but should not expect an ideal-case speedup of 4). Sequential results are from a dedicated sequential implementation, not from a parallel implementation run with a single thread. Timing results include preprocessing time and thread startup costs, but not the time taken to read in the graph file and generate random labels.

Standard Benchmark Problems In Table 1 we present results from the same set of benchmark instances as Carrabs et al. [1]. These are some of the smaller graphs from the DIMACS implementation challenge[1], with randomly allocated labels. Carrabs et al. used three samples for each measurement, and presented the average; we use one hundred. Note that our CPU is newer than that of Carrabs et al., and we have not attempted to scale their results for a "fair" comparison.

The most significant result is that none of our parallel runtime averages are above seven seconds, and none of our sequential runtime averages are above twenty four seconds (our worst sequential runtime from any instance is 32.3 seconds, and our worst parallel runtime is 8.4 seconds). This is in stark contrast to Carrabs et al., who aborted some of their runs on these instances after three hours. Most strikingly, the keller4 instances, which all took Carrabs et al. at least an hour, took under 0.1 seconds for our parallel algorithm. We are using a different model CPU, so results are not directly comparable, but we strongly doubt that hardware differences could contribute to more than one order of magnitude improvement in the runtimes.

We also see that parallelism is in general useful, and is never a penalty, even with very low runtimes. We see a speedup of between 3 and 4 on the non-trivial instances. This is despite the initial sequential portion of the algorithm, the cost of launching the threads, the general complications involved in parallel branch and bound, and the hardware providing only two "real" cores.

Large Sparse Graphs In Table 2 we present results using the Erdős collaboration graphs from the Pajek dataset by Vladimir Batagelj and Andrej Mrvar[2]. These are large, sparse graphs, with up to 7,000 vertices (representing authors) and 12,000 edges (representing collaborations). We have chosen these datasets because of the potential "social network analysis" application suggested by Carrabs et al., where edge labels represent a particular kind of common interest, and we are looking for a clique using only a small number of interests.

For each instance we use 3, 4 and 5 labels, with a budget of 2, 3 and 4. The "3 labels, budget 4" cases are omitted, but we include the "3 labels, budget 3" and "4 labels, budget 4" cases—although the clique sizes are the same (and are equal to the size of a maximum unlabelled clique), we see in a few instances the costs do differ where the budget is 4. Again, we use randomly allocated labels and a sample size of 100.

[1] `http://dimacs.rutgers.edu/Challenges/`

[2] `http://vlado.fmf.uni-lj.si/pub/networks/data/`

Table 2 Experimental results on Erdős collaboration graphs. For each instance, we use three different label set sizes and three different budgets, with randomly allocated labels, and show averages over 100 runs. In each case, we show the average size and cost of the result, and the sequential runtime in seconds.

		budget = 2			budget = 3			budget = 4		
Instance	$\|L\|$	size	cost	t_{seq}	size	cost	t_{seq}	size	cost	t_{seq}
Erdos971	3	5.20	1.98	0.00	7.00	3.00	0.00			
472 vertices,	4	4.61	1.86	0.00	5.71	2.78	0.00	7.00	3.98	0.00
1314 edges	5	4.24	1.94	0.00	5.10	2.81	0.00	6.05	3.91	0.00
Erdos972	3	5.30	1.98	0.12	7.00	3.00	0.12			
5488 vertices,	4	4.84	1.90	0.12	5.84	2.88	0.12	7.00	3.97	0.12
8972 edges	5	4.35	1.80	0.12	5.18	2.76	0.12	6.04	3.82	0.12
Erdos981	3	5.18	1.99	0.00	7.00	3.00	0.00			
485 vertices,	4	4.68	1.91	0.00	5.78	2.88	0.00	7.00	3.99	0.00
1381 edges	5	4.18	1.82	0.00	5.21	2.89	0.00	6.10	3.82	0.00
Erdos982	3	5.29	1.99	0.14	7.00	3.00	0.14			
5822 vertices,	4	4.80	1.85	0.14	5.95	2.95	0.14	7.00	3.96	0.14
9505 edges	5	4.26	1.69	0.14	5.19	2.82	0.14	6.19	3.88	0.14
Erdos991	3	5.18	1.97	0.00	7.00	3.00	0.00			
492 vertices,	4	4.64	1.89	0.00	5.84	2.91	0.00	7.00	3.98	0.00
1417 edges	5	4.23	1.93	0.00	5.19	2.82	0.00	6.16	3.86	0.00
Erdos992	3	5.36	2.00	0.15	8.00	3.00	0.15			
6100 vertices,	4	4.92	1.92	0.15	6.06	2.98	0.15	8.00	3.99	0.15
9939 edges	5	4.27	1.84	0.15	5.32	2.86	0.15	6.38	3.88	0.15
Erdos02	3	5.86	2.00	0.19	8.00	3.00	0.19			
6927 vertices,	4	5.04	1.99	0.19	6.43	2.98	0.19	8.00	3.99	0.19
11850 edges	5	4.66	1.82	0.19	5.69	2.83	0.19	6.82	3.91	0.19

Despite their size, none of these graphs are at all challenging for our algorithm, with average sequential runtimes all being under 0.2 seconds. However, no benefit at all is gained from parallelism—the runtimes are dominated by the cost of preprocessing and encoding the graph, not the search.

4 Possible Improvements and Variations

We will briefly describe three possible improvements to the algorithm. These have all been implemented and appear to be viable, but for simplicity we do not go into detail on these points. We did not use these improvements for the results in the previous section. We also suggest a variation of the problem.

Resuming where we left off: Rather than doing two full passes, it is possible to start the second pass at the point where the last unseating of the incumbent occurred in the first pass. In the sequential case, this is conceptually simple but messy to implement: viewing the recursive calls to `expand` as a tree, we could store the location whenever the incumbent is unseated. For the second pass, we could then skip portions of the search space "to the left" of this

point. In parallel, this is much trickier: it is no longer the case that when a new incumbent is found, we have necessarily explored every subtree to the left of its position.

Different initial vertex orders: We order vertices by non-increasing degree order at the top of search. Other vertex orderings have been proposed for the maximum clique problem, including a dynamic degree and exdegree ordering [14], and minimum-width based orderings [8,9]. These orderings give small improvements for the harder problem instances when labels are present. However, for the Erdős graphs, dynamic degree and exdegree orderings were a severe penalty—they are more expensive to compute (adding almost a whole second to the runtime), and the search space is too small for this one-time cost to be ignored.

Reordering colour classes: For the maximum clique problem, small but consistent benefits can be had by permuting the colour class list produced by `colourOrder` to place colour classes containing only a single vertex at the end, so that they are selected first [6]. A similar benefit is obtained by doing this here.

A multi-label variation of the problem: In the formulation by Carrabs et al., each edge has exactly one label. What if instead edges may have multiple labels? If taking an edge requires paying for all of its labels, this is just a trivial modification to our algorithm. But if taking an edge requires selecting and paying for only one of its labels, it is not obvious what the best way to handle this would be. One possibility would be to branch on edges as well as on vertices (but only where none of the available edges matches a label which has already been selected).

This modification to the problem could be useful for real-world problems: for Carrabs et al.'s example where labels represent different relationship types in a social network graph, it is plausible that two people could both be members of the same club and be colleagues. Similarly, for the Erdős datasets, we could use labels either for different journals and conferences, or for different topic areas (combinatorics, graph theory, etc.). When looking for a clique of people using only a small number of different relationship types, it would make sense to allow only one of the relationships to count towards the cost. However, we suspect that this change could make the problem substantially more challenging.

5 Conclusion

We saw that our dedicated algorithm was faster than a mathematical programming solution. This is not surprising. However, the extent of the performance difference was unexpected: we were able to solve multiple problems in under a tenth of one second that previously took over an hour, and we never took

more than ten seconds to solve any of Carrabs et al.'s instances. We were also able to work with large sparse graphs without difficulty.

Of course, a more complicated mathematical programming model could close the performance gap. One possible route, which has been successful for the maximum clique problem in a SAT setting [4], would be to treat colour classes as variables rather than vertices. But this would require a pre-processing step, and would lose the "ease of use" benefits of a mathematical programming approach. It is also not obvious how the label constraints would map to this kind of model, since equivalently coloured vertices are no longer equal.

On the other hand, adapting a dedicated maximum clique algorithm for this problem did not require major changes. It is true that these algorithms are non-trivial to implement, but there are at least three implementations with publicly available source code (one in Java [8] and two with multi-threading support in C++ [2,5]). Also of note was that bit- and thread-parallelism, which are key contributors to the raw performance of maximum clique algorithms, were similarly successful in this setting.

A further surprise is that threading is beneficial even with the low runtimes of some problem instances. We had assumed that our parallel runtimes would be noticeably worse for extremely easy instances, but this turned out not to be the case. Although there was no benefit for the Erdős collaboration graphs, which were computationally trivial, for the DIMACS graphs there were clear benefits from parallelism even with sequential runtimes as low as a tenth of a second. For the non-trivial instances, we consistently obtained speedups of between 3 and 4. Even on inexpensive desktop machines, it is worth making use of multiple cores.

References

1. Carrabs, F., Cerulli, R., Dell'Olmo, P.: A mathematical programming approach for the maximum labeled clique problem. Procedia - Social and Behavioral Sciences **108**(0), 69 – 78 (2014). DOI http://dx.doi.org/10.1016/j.sbspro.2013.12.821. URL `http://www.sciencedirect.com/science/article/pii/S187704281305461X`
2. Depolli, M., Konc, J., Rozman, K., Trobec, R., Janežič, D.: Exact parallel maximum clique algorithm for general and protein graphs. Journal of Chemical Information and Modeling **53**(9), 2217–2228 (2013). DOI 10.1021/ci4002525. URL `http://pubs.acs.org/doi/abs/10.1021/ci4002525`
3. Lai, T.H., Sahni, S.: Anomalies in parallel branch-and-bound algorithms. Communications of the ACM **27**(6), 594–602 (1984)
4. Li, C.M., Zhu, Z., Manyà, F., Simon, L.: Minimum satisfiability and its applications. In: Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence - Volume Volume One, IJCAI'11, pp. 605–610. AAAI Press, Palo Alto, CA, USA (2011). DOI 10.5591/978-1-57735-516-8/IJCAI11-108. URL `http://dx.doi.org/10.5591/978-1-57735-516-8/IJCAI11-108`
5. McCreesh, C., Prosser, P.: Multi-threading a state-of-the-art maximum clique algorithm. Algorithms **6**(4), 618–635 (2013). DOI 10.3390/a6040618. URL `http://www.mdpi.com/1999-4893/6/4/618`
6. McCreesh, C., Prosser, P.: Reducing the branching in a branch and bound algorithm for the maximum clique problem. In: Principles and Practice of Constraint Programming, 20th International Conference, CP 2014. Springer, Berlin, Heidelberg (2014)

7. McCreesh, C., Prosser, P.: The shape of the search tree for the maximum clique problem, and the implications for parallel branch and bound. ACM Transactions on Parallel Computing (2014). To appear; preprint as CoRR abs/1401.5921
8. Prosser, P.: Exact algorithms for maximum clique: a computational study. Algorithms **5**(4), 545–587 (2012). DOI 10.3390/a5040545. URL http://www.mdpi.com/1999-4893/5/4/545
9. San Segundo, P., Lopez, A., Batsyn, M.: Initial sorting of vertices in the maximum clique problem reviewed. In: P.M. Pardalos, M.G. Resende, C. Vogiatzis, J.L. Walteros (eds.) Learning and Intelligent Optimization, Lecture Notes in Computer Science, pp. 111–120. Springer International Publishing (2014). DOI 10.1007/978-3-319-09584-4_12. URL http://dx.doi.org/10.1007/978-3-319-09584-4_12
10. San Segundo, P., Matia, F., Rodriguez-Losada, D., Hernando, M.: An improved bit parallel exact maximum clique algorithm. Optimization Letters **7**(3), 467–479 (2013). DOI 10.1007/s11590-011-0431-y. URL http://dx.doi.org/10.1007/s11590-011-0431-y
11. San Segundo, P., Rodríguez-Losada, D., Jiménez, A.: An exact bit-parallel algorithm for the maximum clique problem. Comput. Oper. Res. **38**(2), 571–581 (2011). DOI 10.1016/j.cor.2010.07.019. URL http://dx.doi.org/10.1016/j.cor.2010.07.019
12. Tomita, E., Kameda, T.: An efficient branch-and-bound algorithm for finding a maximum clique with computational experiments. Journal of Global Optimization **37**(1), 95–111 (2007)
13. Tomita, E., Seki, T.: An efficient branch-and-bound algorithm for finding a maximum clique. In: Proceedings of the 4th international conference on Discrete mathematics and theoretical computer science, DMTCS'03, pp. 278–289. Springer-Verlag, Berlin, Heidelberg (2003). URL http://dl.acm.org/citation.cfm?id=1783712.1783736
14. Tomita, E., Sutani, Y., Higashi, T., Takahashi, S., Wakatsuki, M.: A simple and faster branch-and-bound algorithm for finding a maximum clique. In: M. Rahman, S. Fujita (eds.) WALCOM: Algorithms and Computation, *Lecture Notes in Computer Science*, vol. 5942, pp. 191–203. Springer Berlin Heidelberg, Berlin, Heidelberg (2010). DOI 10.1007/978-3-642-11440-3_18. URL http://dx.doi.org/10.1007/978-3-642-11440-3_18
15. Trienekens, H.W.: Parallel branch and bound algorithms. Ph.D. thesis, Erasmus University Rotterdam (1990)